Department of Physics and Astronomy
Heidelberg University

Bachelor Thesis in Physics
submitted by

Michael Paping

born in Solothurn (Switzerland)

2023

Characterization of low-power I/O cell input circuits for chip-to-chip interconnects

This Bachelor Thesis has been carried out by Michael Paping at the
European Institute for Neuromorphic Computing (EINC) in Heidelberg
under the supervision of
Dr. Johannes Schemmel

Abstract

BrainScaleS-2 is a mixed-signal neuromorphic architecture that emulates interconnectable neurons. With the goal of creating large networks a test chip was produced that is used to test low-power, source-synchronous chip-to-chip links. This work characterizes the input circuit of the said chip by measuring the receiver's data requirements with well-defined test waveforms. The signal amplitude and timing qualities that are demanded for a reliable data transmission serve as as reference points for the characterization. The results are then compared to the output data characteristics of the same chip that were specified in an earlier work.

In the end, the highest possible data transmission rate that still satisfies a bit error rate of 10^{-5} , limited by the measurement setup, is evaluated and found to be at 2 Gbit/s. However, the result depicts only an upper limit for the transmission rate as the bit error rate is desired to be even lower and the measurements were conducted only on one chip.

Zusammenfassung

BrainScaleS-2 ist eine neuromorphe Architektur, die mithilfe von analogen und digitalen Datensignalen Neuronen und deren Verknüpfung emuliert. Im Zuge der Vergrößerung der Netzwerke durch die Verschaltung von mehreren Chips muss auch deren Kommunikation über energieeffiziente quellsynchrone Verbindungen erforscht werden. Ein eigens entwickelter Testchip liefert die Möglichkeit, die erforderlichen Ein- und Ausgangsschaltungen zu untersuchen.

Gegenstand dieser Arbeit ist die Charakterisierung des Eingangsschaltkreises mithilfe von wohldefinierten externen Stimuli. Als Anhaltspunkt für die Anforderungen des Empfängers dienen dabei die für eine korrekte Erkennung erforderliche Signalamplitude und dessen Timing. Die Ergebnisse werden anschließend mit der zuvor gemessennen Datenqualität der Ausgangsschaltung verglichen, um die höchstmögliche Übertragungsgeschwindigkeit zu bestimmen, die eine Bitfehlerrate von maximal 10^{-5} erfüllt. Dieser Wert ist durch den Messaufbau beschränkt.

Die Ergebnisse lassen eine Datenrate von 2 Gbit/s möglich erscheinen, die jedoch als eine obere Schranke verstanden werden müssen. Da die letzendliche Fehlerrate deutlich niedriger liegen soll und die Daten nur auf der Vermessung eines einzelnen Chips basieren, kann das Ergebnis nur mit Einschränkungen auf andere Systeme übertragen werden.

Contents

1 Introduction

1.1 Motivation

In the context of the developing BrainScaleS-2 neuromorphic architecture it is desired to interconnect chips to larger networks. In doing so, I/O cells are needed that control the input and output of the chip, operating at high speed and low power to ensure sufficient data transmission rates and energy efficiency. A chip was designed specifically to test the I/O cells regarding chip-to-chip connections.

To identify the highest possible operating frequency that still generates acceptably low bit error rates the receiver's input specifications are measured and compared to the output data characteristics, defined in an earlier work [1].

1.2 Theoretical Foundations

1.2.1 Synchronous Digital Signaling

Voltage mode synchronous digital signalling is conducted with two separate data streams. The data signal is used to encode the digital bits by sending consecutive voltage peaks of different amplitudes. A certain voltage V_{OH} is used for a digital "one" and a second voltage V_{OL}, usually ground, for a "zero". The clock provides a common time scheme to synchronize the transmitter and receiver by sending regularly alternating highs and lows at the same rate as the data signal.

The receiver, timed by the clock, then extracts the data by measuring the respective voltage level. If the voltage exceeds a certain value, called V_{IH}, a "one" is detected, if it falls below V_{IL}, a "zero" is recognized. The fact that the level at which a peak is detected is dependent on the signal history is called a hysteresis. In theory, there is no hysteresis as V_{IH} and V_{IL} overlap at $V_{\mathrm{OH}}/2$, but in practice it creates a voltage range where no peak is reliably detected. The hysteresis imposes higher requirements for the input data but is in parts desired for stability against unwanted voltage fluctuations.

The chip uses double data rate, sampling the data twice in every clock cycle, at the rising and the falling edge (c.f. Figure 1.1).

1.2.2 Timing

For the synchronous on-chip logic to reliably detect the correct data, the signal level needs to be reached for a specific time before the clock triggers the data sampling and to be constant shortly afterward. These time spans are referred to as setup and hold time.

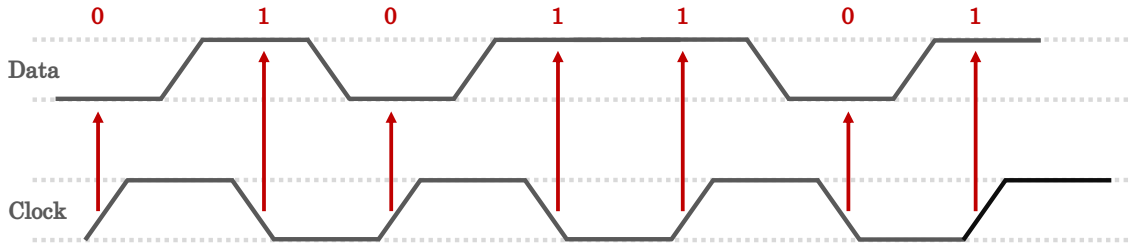

Figure 1.1: Single-ended, voltage mode synchronous digital signaling with double data rate. The source signal is encoded in digital "bits" represented by voltage peaks. A certain voltage V_{OH} implies a "one" and a lower voltage, usually ground, a "zero". A second signal, the clock, consists of steadily alternating "ones" and "zeros", thus providing a synchronization between transmitter and receiver. Double data rate implies that the data is sampled at every rising and falling edge of the clock signal.

The design of the receiver (Figure 1.5) produces the phenomenon that a signal arriving close to the setup or hold time boundaries takes longer time to be detected than an optimal timed signal. The resulting "clock-to-output delay" limits the maximal operating frequency when it amounts close to the value of a clock cycle. In addition, peaks that arrive in these domains are more likely to be falsely detected as the temporal noise, called jitter, potentially makes them violate setup or hold time [2].

1.2.3 Data Mask Test

The receiver's requirements for amplitude and timing can be merged into the so-called input data mask. It imposes a lower limit for the input data shape. Because amplitude and timing are not independent the mask is approximated by a diamond profile. Its height is determined by the distance between V_{IH} and V_{IL} and the width of the minimal signal length that still satisfies the timing constraints.

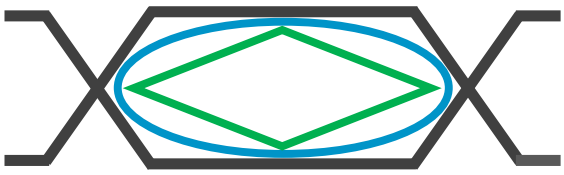

Figure 1.2: Eye pattern of the output data and data mask of the receiver. Overlaying the output data waveforms (grey) creates a pattern that resembles the shape of an eye and depicts a lower limit of the data quality (blue). The receiver's requirements are merged into a diamond-shaped mask (green), representing, in turn, a lower limit for the incoming waveforms. If the mask fits into the eye pattern a reliable data transmission is ensured.

The output data quality of the transmitter can be represented with the so-called data eye. Over a significant amount of time, a series of random data waveforms is sampled and overlayed at each clock trigger, forming an eye-shaped pattern. If the input data mask now fits into this output data eye, reliable data transmission is guaranteed (Figure 1.2).

By designing a possibility to change the delay between clock and data, further referred to as skew or clock-to-data delay, the mask can be adjusted in the time domain. On the other hand, being able to adjust the input reference voltages, the mask position can be optimized vertically.

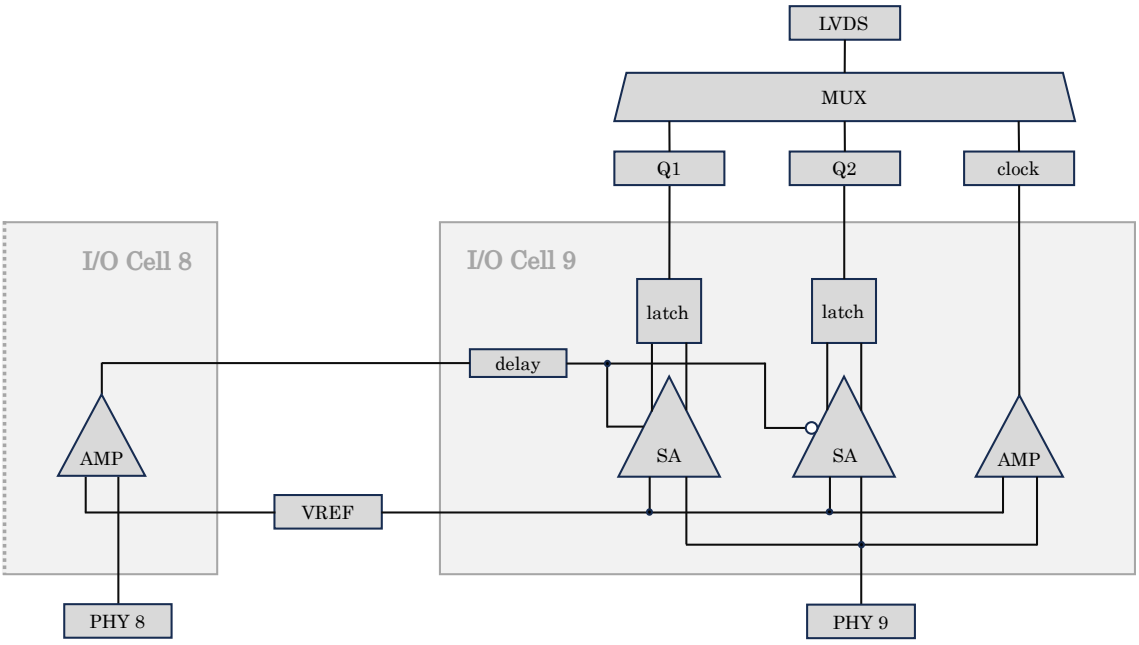

Figure 1.4: I/O cells with LVDS output. The utilized receiver consists of two of the chip's I/O cells, number 8 and 9. One cell has two clock synchronous sense amplifiers and one asynchronous amplifier implemented. They are used for data and clock input, respectively. The clock signal can be sent to adjacent I/O cells to enable the synchronous data reception of the sense amplifiers. They are, in turn, connected to RS latches to form a sense amplifier flip flop (SAFF). As the two SAFFs are timed at opposing clock edges, two data outputs, Q1 and Q2, exist. A delay line enables timing optimization between clock and data. Cell 9 uniquely features a connection to the LVDS debugging port of the chip, being able to switch between three output options, Q1, Q2 and the clock signal.

1.3 The I/O Cell

The measured I/O cells are implemented in a specifically designed test chip (Figure 1.3). It is a 1×1 mm^2 mini@sic manufactured in the 65 nm process. It features a digital control logic and 10 directly accessible high-speed ports (PHYs) from which a single data signal can be transmitted and received. To detect incoming low-power signals ($400 - 1200$ mV) and convert them to the on-chip voltage, the I/O cells are positioned directly at each port (Figure 1.4).

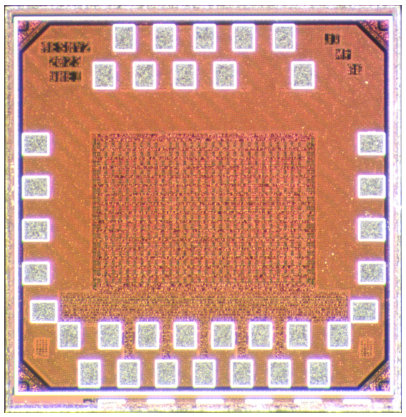

Figure 1.3: Die of the meshy2 test chip, created to asses the I/O cell.

For receiving data, one I/O cell features two synchronous sense amplifiers (Figure 1.5) and one asynchronous amplifier, basically consisting of a differential amplifier and a chain of inverters. The asynchronous amplifier is used to receive the clock which can then be distributed to adjacent I/O cells. It can also be powered down when the I/O cell is used to receive a data signal via the synchronous sense amplifier. In this case, the clock from a nearby I/O cell is utilized. To tune the timing between clock and data, a delay line is interposed that can delay the incoming clock in discrete steps. To enable double data rate, the two synchronous sense amplifiers are sensitive to complementary clock edges. To form

a sense amplifier flip-flop (SAFF), they are each connected to an RS latch which outputs the respective data signal Q1 or Q2.

As cell 9 uniquely features a connection to an LVDS (Low Voltage Differential Signalling) output, three signals, Q1, Q2 and clock can be transferred to that pad via a multiplexer (Figure 1.4). By picking off the according signals, the reception process can be investigated.

The Sense Amplifier Flip-Flop consists of multiple field-effect transistors (FETs) as shown in Figure 1.5. The design is based on Nikolic et al. [3] but inverted to a PFET input stage.

As long as the clock signal is high, P1 remains off and prevents any current flow from the voltage source operating at V_{DDQH}. At the same time, N1 and N4 are active and pull the latch ports R and S down to ground. At the falling clock edge, P1 is turned on, allowing current flow to the two branches. If V_{In} is higher than V_{Ref}, a stronger current flows through the right branch as P2 conducts stronger than P3, pulling the potential of port S higher than port R. Simultaneously, the cross-coupled inverter P4-N2 pulls port R even lower, in turn, pulling port S up again through P5-N3.

These cross-coupled inverters let the circuit settle in one state, in this case, high R and low S, prohibiting any more current flow. This "one" state sets the RS latch. In the other case, the process is inverted, leading to low R and high S which resets the RS latch. The time in which the circuit settles to a state, thus, is longer when the "decision" is tight, as the potentials and currents change more slowly. To operate the SAFF at the rising edge, the clock signal is inverted, deactivating the circuit at the original low level and triggering it at rising clock amplitude.

While the incoming data is fed to V_{In}, V_{Ref} is internally generated and can be adjusted in discrete steps to fit the input data best (c.f. Section 1.2.3).

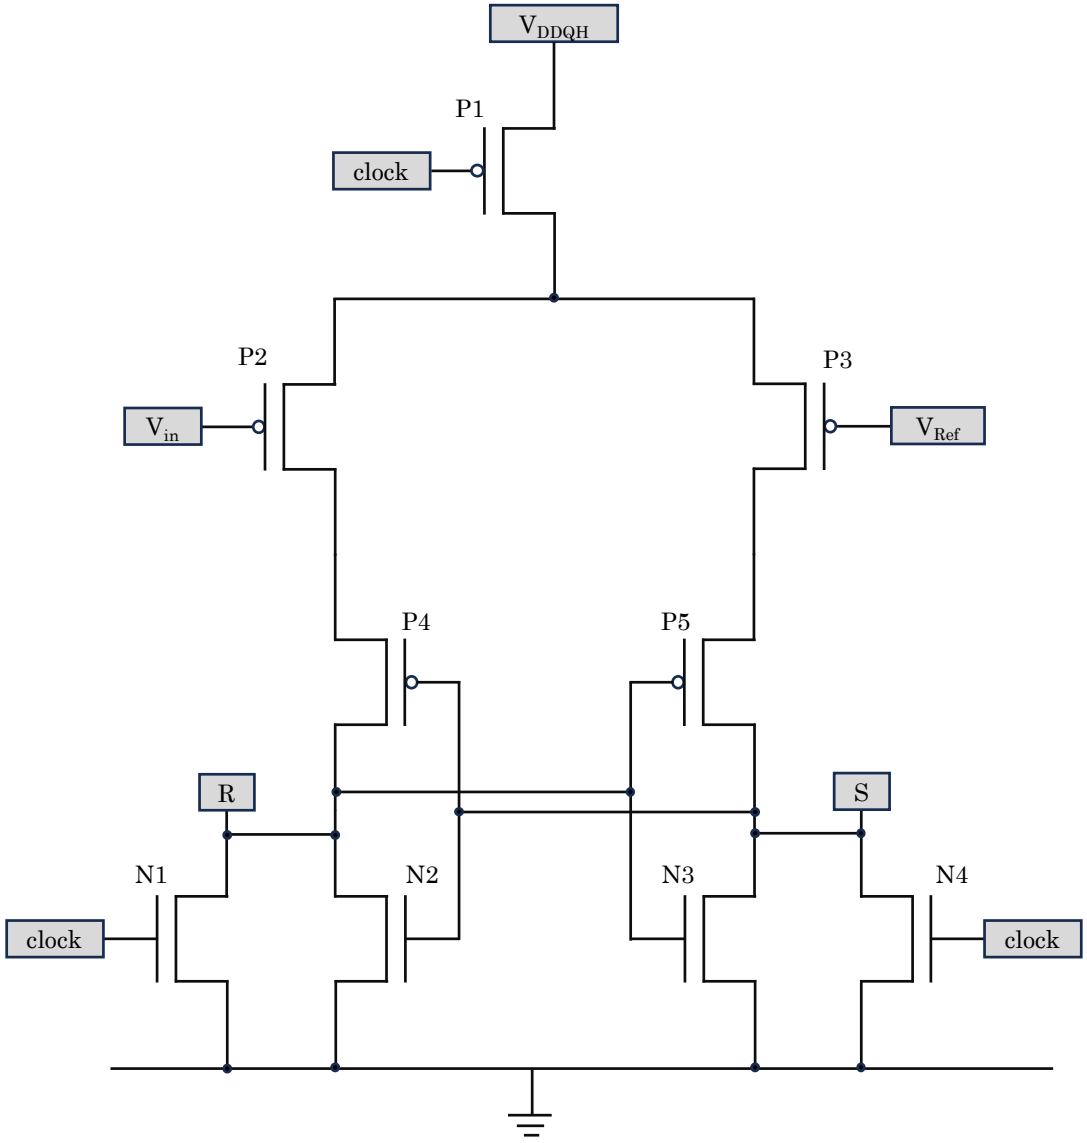

Figure 1.5: The sense amplifier. P1 gates the whole circuit to be inactive at a high clock level. Meanwhile, N1 and N4 are active and completely discharge the ports R and S to ground voltage. As soon as the falling clock edge appears at the circuit, they are turned off while P1 starts to conduct, supplying current to the two branches. Depending on whether V_{In} or V_{Ref} is in a lower state, P2 or P3 increases the current in the respective branch, pulling the potential of port R or S up. The cross coupled inverters (P4-N2, P5-N3) amplify this potential difference until one port is at maximum an the other at minimum potential. The RS latch is thereby set or reset. The circuit stays in this state up to the next clock edge where the ports are discharged again.

2 Experimental Setup

The basic idea of the measurement is to send waveforms of different amplitude and timing to the chip and evaluate the output. Therefore, an arbitrary waveform generator is required to create the particular data and an oscilloscope to analyze the respective output. Additionally, a multimeter is used for monitoring relevant DC parameter values. Remote control of the setup is important to allow long-running measurements. As the instruments are used to the edges of their specifications they are characterized in the following.

2.1 Measurement Equipment

2.1.1 The *Tektronix AWG7102 Arbitrary Waveform Generator*

To provide the Chip with the test waveforms, the *Tektronix AWG 7102* is used (Figure 2.1). It features two analog output channels, each consisting of two single-ended outputs where one of them is inverted. Both of them have a maximal amplitude of 500 mV and a resolution of 2 mV considering 50 Ω termination. This configuration potentially enables a differential operation with an amplitude of 1 V peak-to-peak. The 10 GS/s allow a temporal resolution of 100 ps for the generated waveforms, while a skew of ± 100 ps in steps of 1 ps between the two outputs is also adjustable.

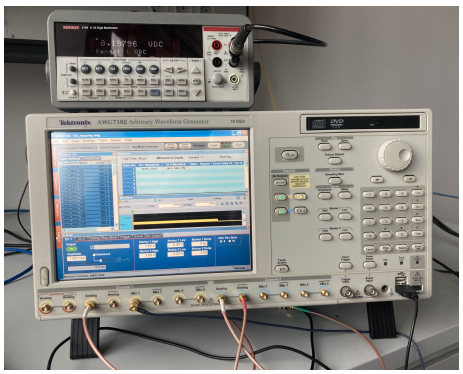

Figure 2.1: The *Tektronix AWG7102 Arbitrary Waveform Generator* and the *Keithley 2100 Multimeter* on top of it.

The differential output is accessible through two SMA jacks, one for the positive part and one for the inverted signal. Another SMA connection supplies a marker signal that is sent out at the beginning of each waveform and can be used as a trigger for the measurement. It is also possible to feed an external clock signal to the AWG to synchronize it to other devices.

The *Tektronix AWG 7102* runs on Windows XP and can be controlled either via a graphical interface and control keys or via GPIB (General Purpose Interface Bus) through LAN.

2.1.2 The *Teledyne LeCroy LabMaster 10-36Zi-A* Oscilloscope

The measurements are conducted with the *Teledyne LeCroy LabMaster 10-36Zi-A* (Figure 2.2). It offers a sampling rate of 80 GS/s when using only two of the four input channels that are accessible through the *LeCroy L2.92A-PLINK Probe Adapter*, again connected to a *LeCroy LPA-BNC* adapter. The oscilloscope offers various tools to extract data from the measured signals like peak-counting, inter-channel delay measurements, high and low-level detection and automatic data saving.

Figure 2.2: The *Teledyne LeCroy LabMaster 10-36Zi-A* Oscilloscope.

2.1.3 The *Keithley 2100 Multimeter*

For the purpose of monitoring the value of V_{Ref} externally, the *Keithley 2100 Multimeter* (Figure 2.1) is connected to the corresponding pad on the PCB. The multimeter is furthermore used in the process of the AWG amplitude characterization (c.f. Section 3.2.3).

2.2 Hardware Modifications

In order to access the chip's I/O cells, two 50 Ω lemo cables are cut and soldered to PHY 8 and 9 of the Chip. As this connection is rather fragile, hot melt adhesive is used to fix the cables on the PCB. Additionally, to measure the LVDS output the *LeCroy WL600 active probe* is soldered to the respective LVDS output pad of the chip and securely fixed with suitable tape (Figure 2.3).

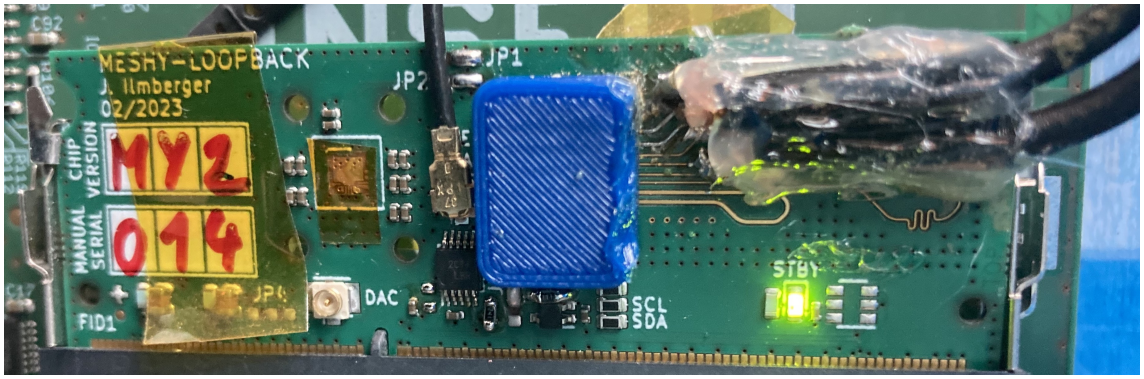

Figure 2.3: The "meshy-loopback" chip carrier. The chip and its I/O cells are placed under the blue cover. On the right, the two lemo cables are soldered to the PHY connections and fixed with hot melt adhesive. Left of the blue cover, the V_{Ref} pad is connected to the multimeter via a mini-coax connection. Not seen on the reverse side is the soldered active differential probe.

2.3 Measurement Setup

The two lemo cables are extended via SMA connectors and plugged into the positive analog outputs of the AWG as clock and data input. The active probe, in turn, is connected via the respective active probe adapter to Channel 1 of the oscilloscope. Channel 3 is connected to an inverted channel output of the AWG to provide a clock or data reference for the measurement. Additionally, the marker output of channel 1 or 3 is linked to the auxiliary output of the oscilloscope. It signals the start of a waveform and is used for trigger purposes. To avoid desynchronization between the two instruments that could distort the timing measurements, the internal clock reference is fed from the oscilloscope to the AWG. This is especially important as the duration of some measurements reaches dozens of hours. At last, all instrument groundings are interconnected via a thick cable to enable compensation currents and establish a common ground (Figure 2.4). An exact overview of the used cables is shown in Table 2.1 and is relevant due to the single-ended nature of the measurement.

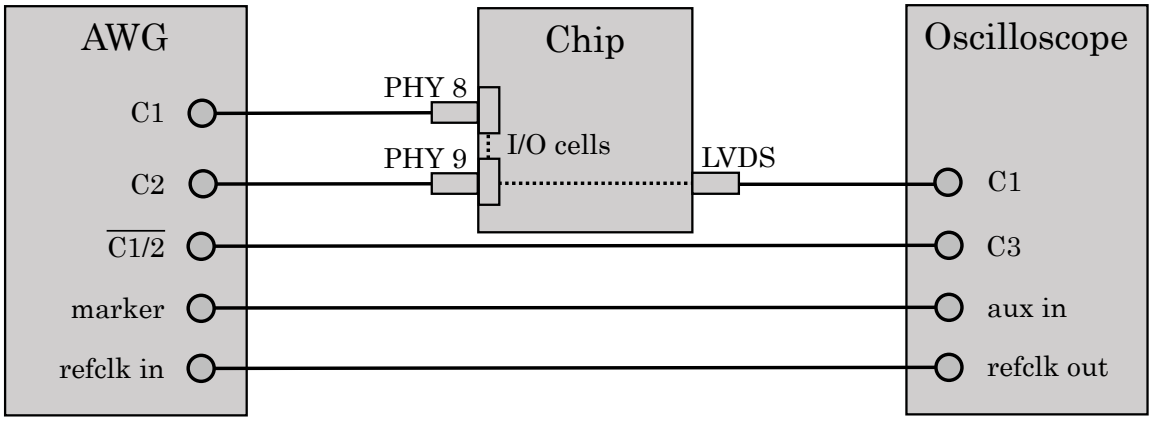

Figure 2.4: This schematic of the measurement setup shows the connections between AWG, chip and oscilloscope. Channel 1 and 2 of the AWG provide the clock and data signals for the chip while one of the inverted data outputs, used for clock or data reference, is connected to the oscilloscope's channel 3. An AWG marker signal that indicates the begin of a waveform is used as a trigger and connected to the auxillary output at the oscilloscope, that, in turn, provides the AWG with an internal clock reference. It is needed to ensure that both instruments work at the same time base even after multiple hours of measurement. The LVDS output of the chip is picked off by the active differential probe and fed to channel 1 of the oscilloscope.

Connection	Cable 1		Cable 2	
	Type	length	Type	length
AWG C1 - PHY8	SMA	50 cm	lemo	82 cm
AWG C2 - PHY9	SMA	50 cm	lemo	80 cm
AWG $\overline{\text{C1/2}}$ - Osci C3	SMA	100 cm	BNC	67 cm
AWG Trigger - Osci Aux	SMA	53 cm	LeCroy Probe Adapter	
LVDS Out - Osci C1	LeCroy WL600 active probe		LeCroy Probe Adapter	

Table 2.1: The cable connection of the measurement setup (Figure 2.4). For pragmatic reasons the connections consist of multiple cable types with different lengths and adapters.

2.4 Measurement Execution

The AWG is controlled with the help of the pyvisa python package [4] and the pyvisa_py backend via GPIB. It features a sequential run mode, where a trigger signal is sent out with each sequence. The oscilloscope, in turn, allows automatic data saving at each trigger, enabling unsupervised data acquisition.

The software and FPGA bitfile versions that were used to configure the chip and conduct the measurements are shown in Table 2.2

repository	short git-hash
chip-io	d713197
chip-io-measurements	e887a94
haldls	edc3d1c
fisch	0d87754
halco	8129637
hxfpga	d4a0922

Table 2.2: Relevant software and FPGA bitfile versions used for the conducted measurements.

3 Results

The following measurements were conducted on the test chip MY2-014 at room temperature. As there were few properly equipped chip-carriers and the hardware modification to enable measurements proved extremely delicate, only one chip example could be probed. As well, the following analysis restraints to only one clock edge.

3.1 Timing Measurements

3.1.1 Introduction

In the course of maximizing the clock frequency and thus the data transmission speed, several limitations become relevant. With rising frequency the signal peaks shorten, still having to satisfy setup and hold time of the receiver. The sum of both gives a first lower boundary for the length of one clock cycle. In addition, the clock-to-output delay starts to increase, leading to another limit for the operating frequency. To get to the limit as close as possible, the timing between clock and data needs to be optimally tuned.

However, the measurement setup prohibits the exact reconstruction of the timing along the transmission paths. This is due to the fact that the according lengths can't be measured sufficiently precise as the 1 ps skew step size of the AWG corresponds to less than 1 mm of cable length. Thus, the skew between clock and data is known at the output of the AWG but shifted by an unknown offset throughout its path to the receiver. In the end, only a relative skew setting is extracted that leads to the best timing configuration.

The goal of this measurement is to find this optimal relative timing between clock and data to ensure the best setup for later measurements. As those results have only setup specific relevance, they are not further discussed.

Secondly, the clock-to-output delay is recorded to estimate a lower limit for the clock cycle length. Again, the recorded delay is shifted by the difference between the two transmission paths of the reference clock and the data signal. As the effective delay can not be determined with reasonable accuracy, it is not possible to extract a minimal pulse width that only allows clock-to-output delays up to an absolute limit. Instead, a pulse width is estimated that avoids relatively high clock-to-output times.

3.1.2 Measurement

As illustrated in Figure 2.4, the AWG sends the clock and data signal to the receiver. The sense amplifier flip-flop detects a “one” or a “zero” and sends the according signal to the oscilloscope via the LVDS output. In order to find a reference point for the timing between clock and data, multiple data waveforms are created that are successively delayed by one AWG data point referring to 100 ps at 10 GS/s sampling rate (Figure 3.1). Combined with the in-built inter-channel skew of the AWG, an overall relative skew range of 1 ns with a resolution of 1 ps is recorded.

By counting the detected peaks at the oscilloscope, a skew domain can be located where the signal is reliably detected. In its center lies, presumably, the best skew setting.

In order to record the clock-to-output delay, the inverted clock output of the AWG is fed to the oscilloscope and the time delta between this reference and the detected LVDS signal is measured. The test waveforms consist of data peaks that are successively skewed around the transition points of dropping detection rate to cover the setup and hold time domains.

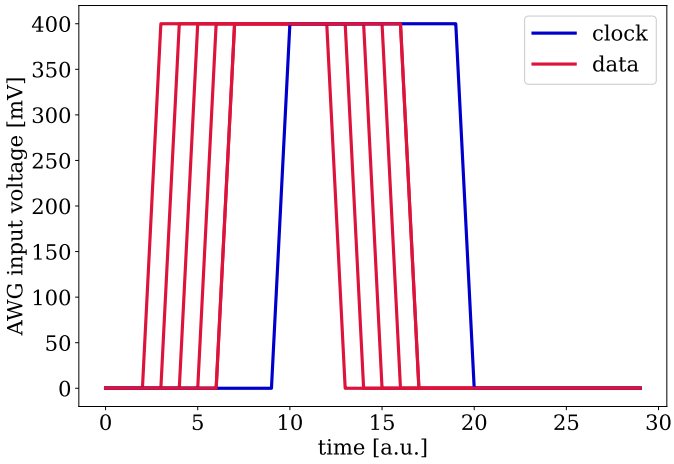

Figure 3.1: Illustration of the AWG input data for the timing measurement (not to scale). As the skew between data and clock signal changes along the transmission path in a manner that is hard to measure, a sweep of gradually delayed peaks is sent to the Chip to determine the best relative skew setting for succeeding measurements regarding peak detection.

To gain a robust statistic one measurement run of all skew settings is then repeated in its entirety several hundred times. This procedure makes sure that any shift in parameters does not lead to a systematic bias, as the measurement duration mounts up to many hours.

3.1.3 Results

Firstly, the observation is made that the repeated clock-to-output measurements of one skew setting do not follow a normal distribution and are widely spread over a range of about 300 ps (Figure 3.2). The expected sources of noise in the time domain, known as jitter, are the clock and data output of the AWG, all on-chip circuits, especially the LVDS output and the active differential probe.

A jitter analysis shows that the AWG’s output timing only changes about a few ps while the signal from the active probe deviates up to 50 ps around the mean value. However, this is only true for the timing of the active probe’s edges, while the width of the peaks change in a similar manner as the recorded clock-to-output times, spreading in a range of about 300 ps. This bias likely is the cause for the observed data distribution and is believed to originate from the active probe.

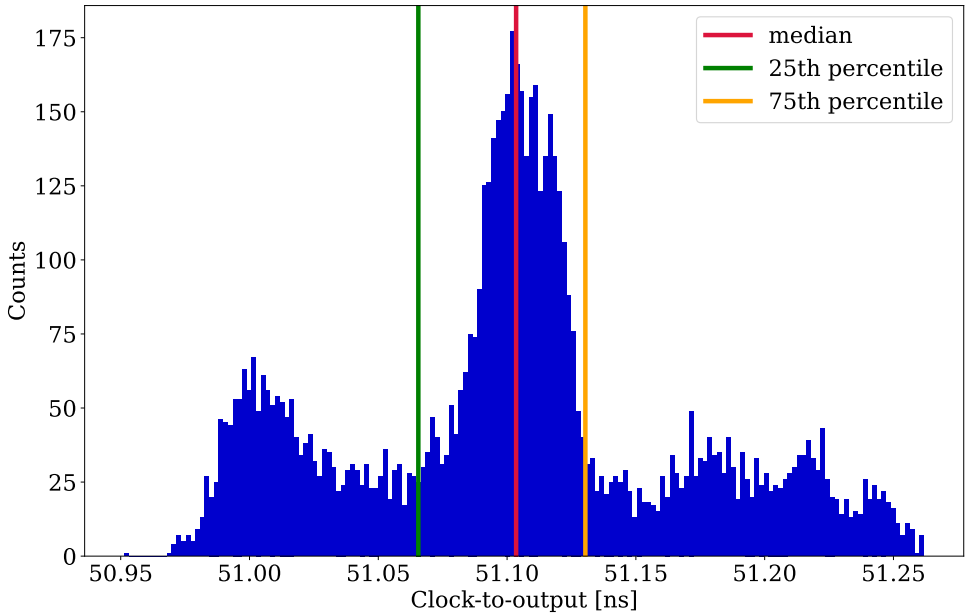

Figure 3.2: The clock-to-output time distribution for one skew setting. The multiple jitter sources, namely the AWG output, the on-chip circuits with the LVDS output and the active probe, should lead to a normal distribution of the clock-to-output times. However, the measured distribution is asymmetric, consists of several peaks and spreads along about 300 ps, a multiple of the individual jitter sources. A precise analysis of the ouput from the active probe shows that the peaks arrive with a very small temporal deviation but highly differ in width. This bias is likely to be an explanation for the recorded data distribution and is believed to be caused by the active probe. Thus, the extreme jitter is not considered as influential for the results, while the data has to be treated as arbitrarily distributed regardless.

The recorded clock-to-output times as a function of the inter-channel skew between clock and data are shown in Figure 3.3. While all the measured time differences are affected by the discussed offset, a clear rise of the delay is seen at the transition point mounting up to about 80 ps considering the respective medians. The data suggests that a minimum data peak width of about 85 ps should be maintained to avoid the domain of high delays at optimal timing. This value is, however, additionally enlarged by the imperfection of the timing adjustment.

3.2 Amplitude Measurements

With the optimal timing settings from the previous measurements, which were conducted with the maximal amplitude, now the amplitude measurements are executed. This reflects the approach of an input mask based on four edge data points: amplitude range at perfect timing and minimal peak width with perfect amplitude.

3.2.1 Introduction

The input data has a set maximum amplitude of $V_{\mathrm{IH}} = 400\,\mathrm{mV}$. Therefore, V_{Ref} is initially adjusted to about 200 mV. The actual voltage levels that define the ranges

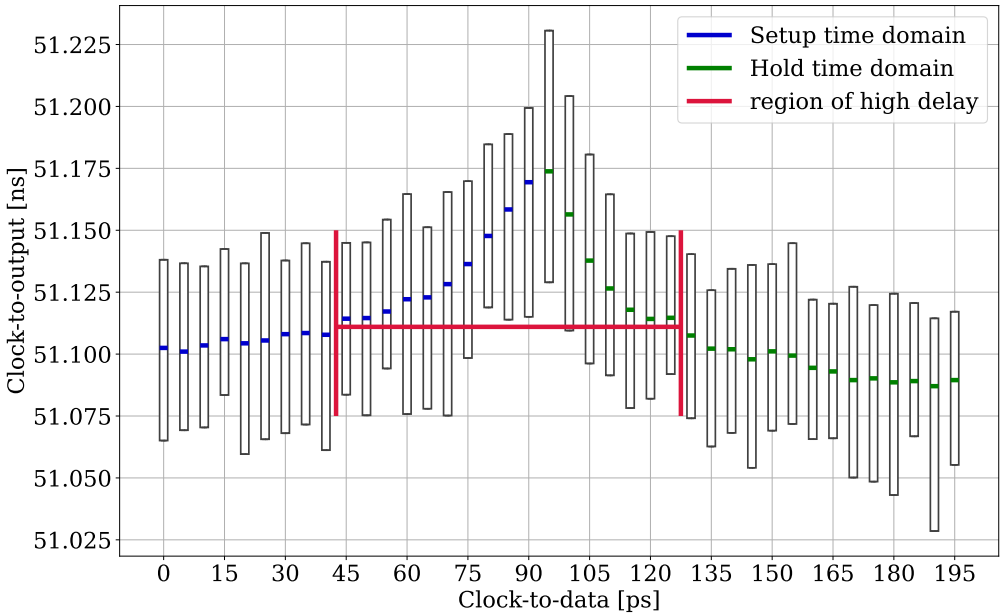

Figure 3.3: The clock-to-output delay measurement. The x-axis depicts the inter channel skew between clock and data signal, measured at the AWG. The y-axis represents the transmission time of one data peak from AWG to the oscilloscope. To give consideration to the arbitrary data distribution, the colored patches indicate the median while the upper and lower boxes mark the adjacent quartiles. The two branches represent the hold and setup time domains. As the skew progresses further, the time that the receiver needs to detect the signal state gets increasingly longer until it is no longer detected reliably at all due to setup or hold time violation. Because of the cable setup and transmission paths through the chip, both axes are shifted by an unknown offset. As the retrieval of an exact absolute time scheme is barely possible with the measurement setup only the relative timing quantities are used. As significantly delayed signals can interfere with subsequent ones, a minimal pulse width is defined that prohibits such a delay at optimal timing. In this case it is estimated to be 85 ps, not taking into account the delay tuning constraints.

of reliable signal detection, V_{IH} and V_{IL}, are by design expected to be located around it. This hysteresis is helpful to avoid false detection due to oscillations in the signal but enlarges the input data mask. As the sense amplifier consist of FETs that are operated at V_{DDQH}, it's variation is believed to have an significant impact on the location of the characteristic amplitude levels.

The asynchronous receiver, effectively being a differential amplifier and a chain of inverters, receives only one signal and decides continuously in time, whether the voltage level for a peak is reached. In this case, a hysteresis is even more important as signal fluctuations can lead to false detection at any point in time.

The aim of the measurement is hence to characterize both the synchronous and asynchronous amplifier regarding the location of V_{IH} and V_{IL} and to investigate the influence of a fluctuating operating voltage V_{DDQH}.

3.2.2 Measurement

In order to do so, clock and data are again fed to the receiver as shown in Figure 2.4. The data consists of pulses with different amplitudes, using the previously

determined optimal timing (Figure 3.4). The output of the sense amplifier flip-flop is fed to the oscilloscope via the LVDS pad (c.f. Figure 2.4). One measurement run consists of a few thousand peaks with equal amplitude where each data peak is environed by 3 clock cycles of constant low or high level, respectively (Figure 3.4). This is done to ensure that the detection rate covers the "worst case" concerning the signal history.

For the characterization of the asynchronous receiver, just one clock signal with various voltage levels is sent to PHY 9 and the according output signal is fed to the oscilloscope.

Several runs with different amplitudes make up one sweep that is, in turn, repeated to gain statistics. This procedure makes sure that any drift of parameters in time does not result in a systematic bias of the measurement. In the end, the oscilloscope counts the detected peaks per run, from which the detection rate can be computed. To investigate the influence of V_{DDQH} on the results, it is varied by $\pm\,10\,\%$ around the default value of 1.2 V.

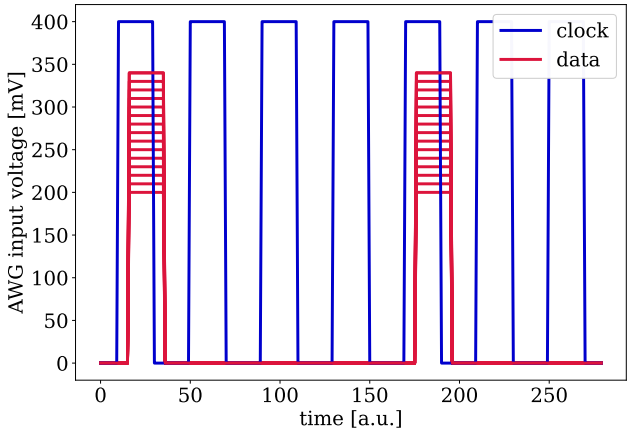

Figure 3.4: Illustration of the ideal input data for the AWG, measuring "ones" in an "zero" environment (not to scale). Digital waveforms of different heights are fed to the AWG, which then sends the according analog signal to the chip. Every four clock cycles a peak of a certain amplitude is generated to assume the "worst" signal history regarding peak detection. While the amplitude stays constant over one waveform, it is swept over the desired range in consequent runs. The depicted waveforms show the actual timing between clock and data as sent yb the AWG the different transmission lines resulting in an optimal timing at the receiver (c.f. Section 3.1)

3.2.3 AWG Amplitude Characterization

Because the AWG is working at its maximum specifications the rise time prohibits a correct signal amplitude (c.f. Figure 3.5a). This means that in one clock cycle the AWG, switching from low to high, actually produces a signal with a lower amplitude than expected and vice versa.

To correct this deviation, the input data is also fed to the oscilloscope and its real level is measured. As the deviation seems to be linear, a simple fit is made to later estimate what the real amplitude values are (c.f. Figure 3.5b). As the measurement setup allows amplitudes from 0 to 500 mV the actual range of 400 mV is covered in the end.

3.2.4 Results

The results of the synchronous amplifier are shown in Figure 3.6. The branch with the point markers corresponds to the detection rate of "zeros" in a "one" environment, the branch with the triangles to "ones" in a "zero" environment. In both cases the rate is shaped sigmoid-like in the domain of unreliable detection,

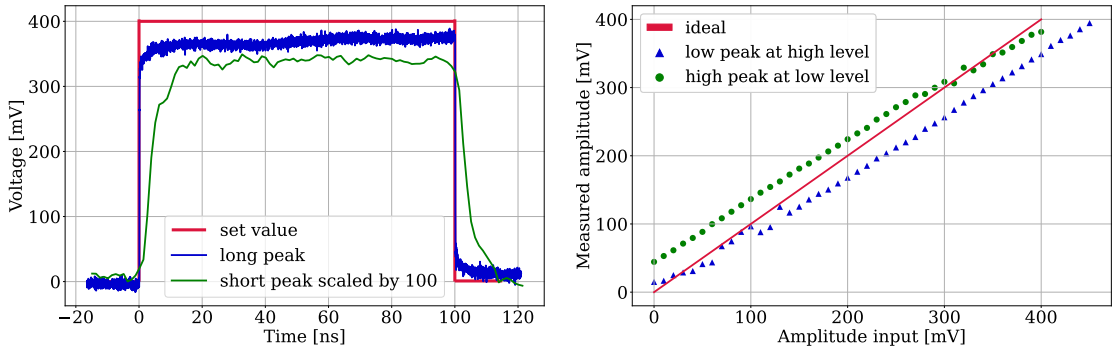

(a) (b)

Figure 3.5: Visualization of the actual AWG output waveform (a). It is observed that the output amplitudes do not reach the set value in time at 500 MHz clock frequency due to the long rise and fall times. The red line indicates the desired value of 400 mV, the blue waveform is a measurement of a peak with a length of 100 ns, the green waveform is a peak with a length of 1 ns scaled by a factor of 100 in time. The slow approach to the set value clearly demonstrates the need of an amplitude correction. AWG Amplitude characterization (b): The plot shows an approximately linear deviation of the real amplitude levels from the desired ones. The actual voltage ranges in the later plots are adjusted by a linear fit. The outliers that do not follow the linear deviation are believed to be fragments of the oscilloscope measurement and thus do not influence the actual results.

together forming the hysteresis.

The plots also show that the dimension and position of this domain is highly dependent on the value of V_{DDQH}. Increasing it corresponds to a stronger hysteresis, decreasing it leads to a weaker one. While the varying location of the domain can in part be explained by a change in V_{Ref}, the configuration with $V_{\mathrm{DDQH}} = 1.32\,\mathrm{V}$ stands out and can not be explained.

In general, it has to be kept in mind that the detection rate for one amplitude level is altogether measured about 10^5 times while the actually desired certainty would be reached at about 10^{10} samples. This means that the actual distance between V_{IL} and V_{IH} could be even greater.

The detection rate of the asynchronous receiver (Figure 3.6d) shows a similar behavior as the synchronous one. Overall, the curve is shifted to higher voltages with a relatively large domain, where the detection rate is close to 100 %, which can not be explained.

Contrary to the synchronous receiver, the variation of V_{DDQH} has no significant impact on the shape of the detection rate curve. However, it is believed to be strongly dependent on the bias current in the amplifier, while its influence on the result was not researched in this work.

3.3 The Data Mask

3.3.1 Introduction

From the level and timing measurements, the height and width of the input data mask can be calculated. The voltage level V_{IH} where a “one” is reliably detected in

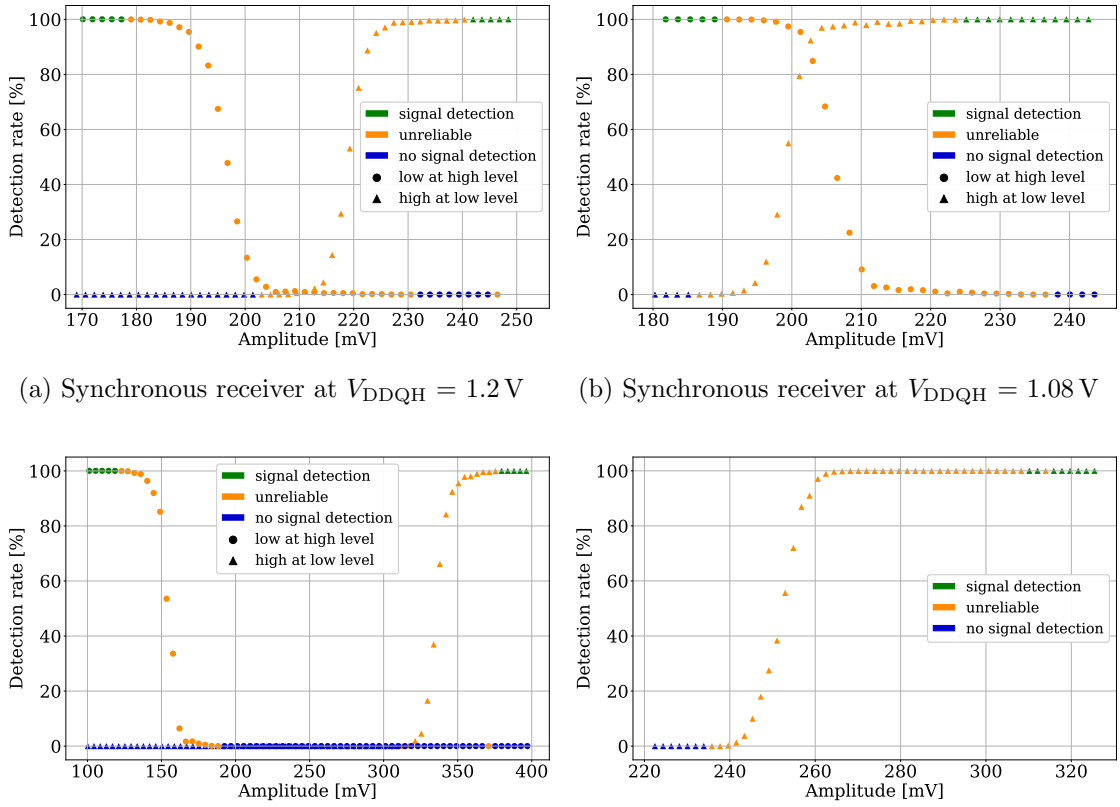

(a) Synchronous receiver at $V_{\mathrm{DDQH}} = 1.2\,\mathrm{V}$

(b) Synchronous receiver at $V_{\mathrm{DDQH}} = 1.08\,\mathrm{V}$

(c) Synchronous receiver at $V_{\mathrm{DDQH}} = 1.32\,\mathrm{V}$

(d) Asynchronous receiver at $V_{\mathrm{DDQH}} = 1.2\,\mathrm{V}$

Figure 3.6: Detection rates of the receiver as a function of input data amplitude. In order to find the voltage level at which a signal is reliably detected by the synchronous receiver, peaks of different heights are sent to the synchronous receiver accompanied by a clock. The detected signals are then sent to the oscilloscope, which counts the peaks and calculates the according rate. The two branches reflect the detection of a "one" in a "zero" environment and the other way round. The detection rate at the default V_{DDQH} level is shown in (a). The detection rates for variations of it are shown in (b) and (c). The domain of unreliable signal detection clearly is highly dependent on V_{DDQH}. For lower voltages the hysteresis diminishes while it enlarges for higher ones. Its position changes as well, which is in part explained by the simultaneously changing V_{Ref}. It is generated internally and also dependent on V_{DDQH}. However, this can not explain the major shift in the 1.32 V case. For the asynchronous receiver (d) just one clock signal with varying amplitude is sent to PHY 9 of the chip. The detected clock peaks are then sent to the oscilloscope through the LVDS output where the detection rate is calculated. Compared to the synchronous receiver the point of reliable signal detection is shifted to higher voltages with a long domain of a detection rate close to 100 % prior to it. While a change of V_{DDQH} has a marginal effect on the shape of the curve, a variation of the bias current is likely to have a strong impact. However, this parameter was set to a standard value and its variation not investigated within the measurement campaign.

a "zero" environment delivers the upper boundary of the mask, whereas the lower boundary is obtained from the "zero" detection threshold V_{IL} (c.f. Figure 3.6). Reliably means that the detection rate is 100 % when sampling, in this case, 10^5 times.

The width is calculated from multiple sources. Firstly, the clock-to-output diagram is considered (c.f. Figure 3.3): To avoid the high delay at the transition points the signal needs to cover at least this domain. Additionally, the discrete delay line steps need to be taken into account. In the worst scenario, the optimally timed point is in between two steps, leading to an additional required signal length of half of the delay line's step size. As the jitter of the clock and data outputs should be covered in the eye measurement, it is not considered here.

These dimensions are then transformed into a diamond-shaped mask that represents the receiver requirements. The size of the input mask is thereby fixed. However, the position is adjustable as V_{Ref} is adjustable and the delay line allows temporal tuning. Its discrete steps are already taken into account in the width calculation. A comparison with the output driver data eye shows how well the transmitter and receiver match at various operating frequencies.

3.3.2 Input Data Mask Calculation

A Breakdown of the amplitude measurements yields very different mask height results depending on V_{DDQH}. The upper boundary ranges from 220 mV to 375 mV while the lower boundary reaches values between 125 mV to 190 mV. The time span for the setup and hold time requirements amounts to about 85 ps (c.f. Figure 3.3), the delay line steps are discretized in steps with a maximum size of 170 ps. This value is dependent on V_{DDQH} as well, however, it is not further analyzed but used as a worst-case scenario. This leads to an overall mask width of 170 ps. The results are summarized in Table 3.1.

3.3.3 Comparison with the Output Driver

In an earlier work [1] the output data eye was measured with the same Teledyne LeCroy Oscilloscope (Section 2.1.2) for various clock frequencies. As it was only possible to retrieve screenshots of the eye patterns, the input masks are fitted to the respective scales that show the voltage of a terminated input, i.e. half the amount, and hand-placed into the best position. To cover the best and worst case, the voltage spans of the two extreme V_{DDQH} cases are used for the respective mask heights (Figure 3.7).

V_{ddqh}	height	width
1.08 V	30 mV	170 ps
1.2 V	60 mV	170 ps
1.32 V	250 mV	170 ps

Table 3.1: The input mask dimensions. The height is retrieved from the amplitude measurements and reflects the voltage span in which the signal is not reliably detected. The width is a combination of the minimal signal length that avoids too high clock-to-output delays and the delay tuning that is only possible in discrete steps.

The results show a good mask matching for the 500 MHz and 1 GHz configuration in both cases, as the diamond mask fits well into the data eye. Especially in the best-case scenario, there are large margins both in temporal and voltage dimension. At 1.5 GHz this is no longer the case as the eye size diminishes vastly due to rare but destructive jitter. Overall, the time domain seems to have larger margins than the voltage level, which is mainly due to the large hysteresis

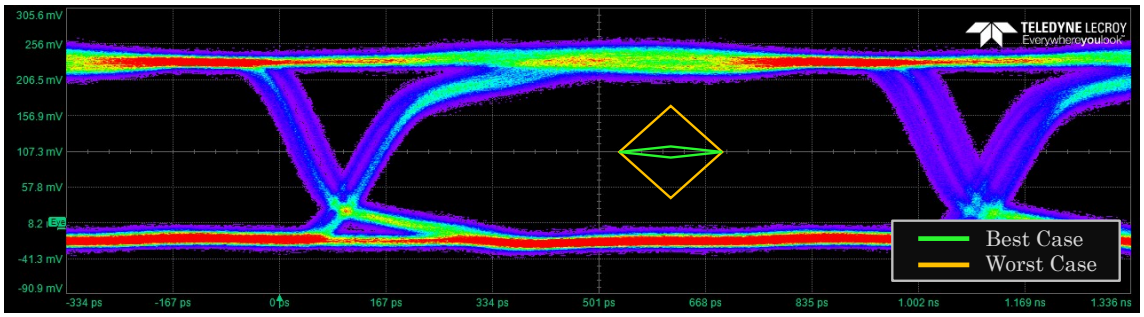

(a) $f_{\mathrm{clk}} = 500\,\mathrm{MHz}$

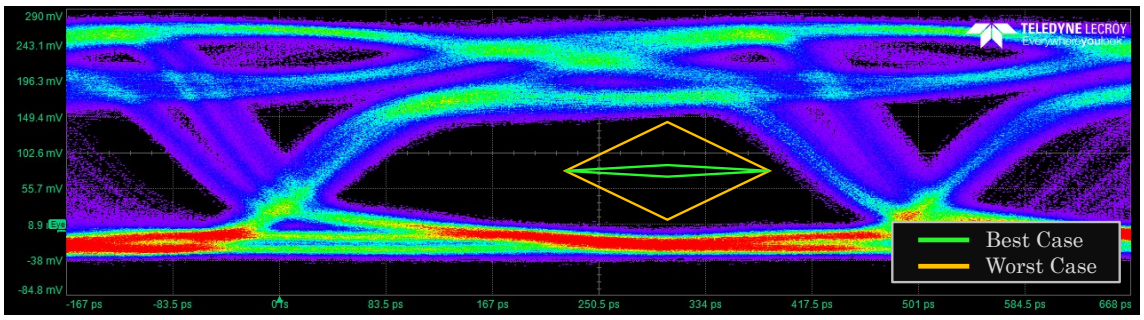

(b) $f_{\mathrm{clk}} = 1\,\mathrm{GHz}$

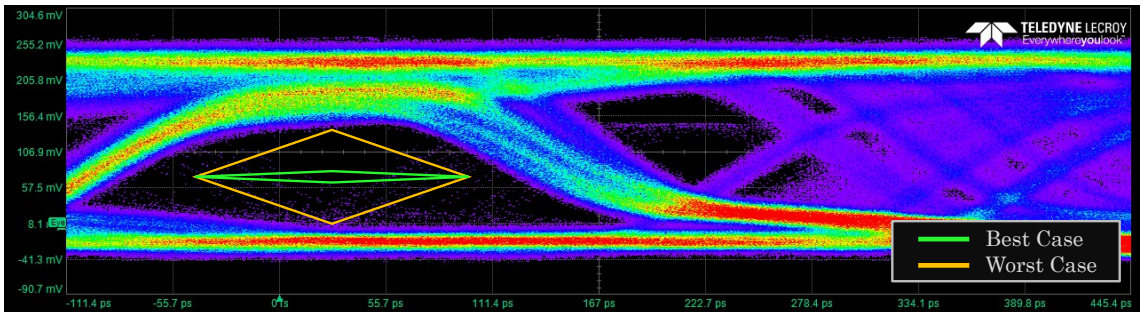

(c) $f_{\mathrm{clk}} = 1.5\,\mathrm{GHz}$

Figure 3.7: Data eye of the output driver with input mask. With the results of the amplitude and timing measurements (Table 3.1) a diamond-shaped mask is defined that represents the receiver's input data requirements. It is then fit into the output data eye from previous measurements to asses for which frequencies the output data quality matches the receiver's specifications.

at high V_{DDQH}.

As the position of the eye is important particularly in the voltage axis, the tuning steps for V_{Ref} need to be considered. As for the best case scenario the vertical placement is not relevant, the positioning of the worst case mask at higher frequencies needs to be rather exact.

The data suggests implementing a precise regulation of V_{Ref} especially below 200 mV to operate at higher frequencies. However, the relatively low statistic has to be kept in mind. Expanding the measurement iterations to values of 10^{10} samples per data point the input mask could be enlarged even further. The result should therefore be treated as reference points for further investigation.

4 Conclusion and Discussion

The timing measurements clearly show the expected clock-to-output delay characteristic in regimes of fading hold and setup time. However, the underlying data samples follow an unexpected and asymmetric distribution. The exact source of this distortion remains unclear but it is believed to be an artefact of the chips output circuit and the active differential probe. Continuing measurements to further isolate the exact source were not conducted but the main contribution is believed to be from the active probe.

Using the respective medians of the distribution as a reference, a minimal pulse width was estimated that avoids too high clock-to-output delays under the premise of perfect timing. Together with the restrictions of the delay line, where a worst-case scenario was assumed, the actual lower limit for a pulse width was specified. Possible influences of a varying V_{DDQH} were not investigated.

The amplitude measurements resulted in voltage thresholds that border regimes of reliable signal detection. A strong dependency of V_{DDQH} was observed that shifted and enlarged the respective hysteresis. However, measurements were handicapped by the AWG limitations as its rise time produces deviating effective peak voltage levels. Therefore, a characterization was conducted to correct the amplitude values. It was performed simultaneously with the data acquisition with the oscilloscope. As its input channels are properly terminated with $50\,\Omega$ whereas the input of the I/O cell is not terminated at all, the results could be biased. In addition, V_{Ref} is generated internally, being also dependent on V_{DDQH}. Although the resulting deviations were relatively small and only the relative dimensions of the hysteresis were extracted, it imposes a small uncertainty. Lastly, strong restrictions to the exact positioning of the characteristic voltage levels have to be made as the statistic is well underneath the desired dimension. The configuration is intended to run with bit error rates of about 10^{-10} while these measurements only support conclusions for rates of about 10^{-5} In theory, an analytical function could be used to fit the data and calculate the voltage levels for the desired accuracy. However, the theoretical basis for such a function was not available.

Both characterizations were then merged into the input data mask and afterwards compared to the output eye pattern. Because of the strongly divergent results in the voltage dimension, a best-case and a worst-case scenario was established. It was found that input mask and eye pattern match for operating frequencies of 500 MHz and 1 GHz while none matches at 1.5 GHz due to a drastically reduced size of the data eye. Because the data transmitter is designed for short chip-to-chip connections, the eye pattern could be influenced by the measurement setup, which featured comparably long transmission lines. In the end, the output data

quality could therefore be better than represented. Furthermore, as many worst-case assumptions were made, the setup is expected to perform better in most cases in the end. However, both the ouput data eye and the input data mask were recorded with fewer statistics than required. Hence, their dimensions could be over- and underestimated, respectively. Also, the mask measurements focused on one clock edge only, neglecting possible deviations caused by the different shapes of rising and falling edges. Another limitation that needs to be accounted for is the fact that only one chip with its I/O cells was tested at all, disregarding manufacturing variations between different chips.

In a nutshell, the I/O cells are expected to work with a frequency in the domain of 1 GHz, corresponding to a data transmission rate of 2 Gbit/s. To validate this result, it is suggested to conduct further measurements. Possible improvements can be made by taking into account both edges of the clock, sampling more data to create a larger statistic, trying to establish an absolute timing reconstruction and using more chips as a database. However, the latter measures are extremely elaborate and the question has to be posed, if they are worth the benefits.

Bibliography

[1] Michael Paping. *Calibration and Measurement of an I/O Cell Output Driver*. Internship report. 2023.

[2] V. L. Chi. "Salphasic Distribution of Clock Signals for Synchronous Systems". In: *IEEE Trans. Comput.* 43.5 (May 1994), pp. 597–602. ISSN: 0018-9340. DOI: 10.1109/12.280806. URL: https://doi.org/10.1109/12.280806.

[3] B. Nikolic et al. "Improved sense-amplifier-based flip-flop: design and measurements". In: *IEEE Journal of Solid-State Circuits* 35.6 (2000), pp. 876–884. DOI: 10.1109/4.845191.

[4] PyVISA. URL: https://github.com/pyvisa/pyvisa.

Danksagung

Zuallererst möchte ich meinen Betreuer Joscha Ilmberger danken, der mit seiner geballten Kompetenz die meisten meiner Denkfehler und Wissenslücken im Handumdrehen in Luft aufgelöst hat, nur um mich dann durch weitergehende Schilderungen zu den Untiefen des digitalen Chipdesigns wieder leicht verwirrt zurückzulassen. Dass er neben seinen Tutortätigkeiten auch noch die Pflege des Tischkickers und der Dartsscheibe zu seinen Pflichten zählt, ist ihm ebenso hoch anzurechnen.

Weiterhin geht mein Dank an Jakob Kaiser für die sorgfältige und gewissenhafte, und an Yannik Stradmann für die pointierte und schonungslose Korrektur dieser Bachelorarbeit, die sonst stark an Qualität eingebüßt hätte.

Jakob Kaiser muss auch in seiner Rolle als Sparringspartner beim Kickern erwähnt werden, die er zusammen mit Tim Auberer mustergültig erfüllt hat. Mit einem 6:0 in der Tasche arbeitet es sich am Nachmittag einfach besser.

Zuletzt gilt mein Dank Bela, Lisa und Anni. Sie haben für diese Arbeit zwar nichts geleistet, aber das sollte auch kein Maßstab für eine Danksagung sein.

Erklärung

Ich versichere, dass ich diese Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den 22.12.2023